

**DATA RECOVERY FOR VIRTUAL ORDERED WRITES
FOR MULTIPLE STORAGE DEVICES**

Background of the Invention

1. Technical Field

5 This application relates to computer storage devices, and more particularly to the field of transferring data between storage devices.

2. Description of Related Art

Host processor systems may store and retrieve data using a storage device containing a plurality of host interface units (host adapters), disk drives, and disk interface units (disk adapters). Such storage devices are provided, for example, by EMC Corporation of Hopkinton, Mass. and disclosed in U.S. Patent No. 5,206,939 to Yanai et al., U.S. Patent No. 5,778,394 to Galtzur et al., U.S. Patent No. 5,845,147 to Vishlitzky et al., and U.S. Patent No. 5,857,208 to Ofek. The host systems access the storage device through a plurality of channels provided therewith. Host systems provide data and access control information through the channels to the storage device and the storage device provides data to the host systems also through the channels. The host systems do not address the disk drives of the storage device directly, but rather, access what appears to the host systems as a plurality of logical disk units. The logical disk units may or may not correspond to the actual disk drives. Allowing multiple host systems to access the 15 single storage device unit allows the host systems to share data stored therein.

20

In some instances, it may be desirable to copy data from one storage device to another. For example, if a host writes data to a first storage device, it may be desirable to copy that data to a second storage device provided in a different location so that if a disaster occurs that renders the first storage device inoperable, the host (or another host) 5 may resume operation using the data of the second storage device. Such a capability is provided, for example, by the Remote Data Facility (RDF) product provided by EMC Corporation of Hopkinton, Massachusetts. With RDF, a first storage device, denoted the "primary storage device" (or "R1") is coupled to the host. One or more other storage devices, called "secondary storage devices" (or "R2") receive copies of the data that is 10 written to the primary storage device by the host. The host interacts directly with the primary storage device, but any data changes made to the primary storage device are automatically provided to the one or more secondary storage devices using RDF. The primary and secondary storage devices may be connected by a data link, such as an ESCON link, a Fibre Channel link, and/or a Gigabit Ethernet link. The RDF 15 functionality may be facilitated with an RDF adapter (RA) provided at each of the storage devices.

RDF allows synchronous data transfer where, after data written from a host to a primary storage device is transferred from the primary storage device to a secondary storage device using RDF, receipt is acknowledged by the secondary storage device to 20 the primary storage device which then provides a write acknowledge back to the host. Thus, in synchronous mode, the host does not receive a write acknowledge from the

primary storage device until the RDF transfer to the secondary storage device has been completed and acknowledged by the secondary storage device.

A drawback to the synchronous RDF system is that the latency of each of the write operations is increased by waiting for the acknowledgement of the RDF transfer.

- 5 This problem is worse when there is a long distance between the primary storage device and the secondary storage device; because of transmission delays, the time delay required for making the RDF transfer and then waiting for an acknowledgement back after the transfer is complete may be unacceptable.

It is also possible to use RDF in an a semi-synchronous mode, in which case the

- 10 data is written from the host to the primary storage device which acknowledges the write immediately and then, at the same time, begins the process of transferring the data to the secondary storage device. Thus, for a single transfer of data, this scheme overcomes some of the disadvantages of using RDF in the synchronous mode. However, for data integrity purposes, the semi-synchronous transfer mode does not allow the primary
- 15 storage device to transfer data to the secondary storage device until a previous transfer is acknowledged by the secondary storage device. Thus, the bottlenecks associated with using RDF in the synchronous mode are simply delayed by one iteration because transfer of a second amount of data cannot occur until transfer of previous data has been acknowledged by the secondary storage device.

Another possibility is to have the host write data to the primary storage device in asynchronous mode and have the primary storage device copy data to the secondary storage device in the background. The background copy involves cycling through each of the tracks of the primary storage device sequentially and, when it is determined that a particular block has been modified since the last time that block was copied, the block is transferred from the primary storage device to the secondary storage device. Although this mechanism may attenuate the latency problem associated with synchronous and semi-synchronous data transfer modes, a difficulty still exists because there can not be a guarantee of data consistency between the primary and secondary storage devices. If there are problems, such as a failure of the primary system, the secondary system may end up with out-of-order changes that make the data unusable.

A proposed solution to this problem is the Symmetrix Automated Replication (SAR) process, which is described in pending U.S. patent applications 10/224,918 and 10/225,021, both of which were filed on August 21, 2002. The SAR uses devices (BCV's) that can mirror standard logical devices. A BCV device can also be split from its standard logical device after being mirrored and can be resynced (i.e., reestablished as a mirror) to the standard logical devices after being split. In addition, a BCV can be remotely mirrored using RDF, in which case the BCV may propagate data changes made thereto (while the BCV is acting as a mirror) to the BCV remote mirror when the BCV is split from the corresponding standard logical device.

However, using the SAR process requires the significant overhead of continuously splitting and resyncing the BCV's. The SAR process also uses host control and management, which relies on the controlling host being operational. In addition, the cycle time for a practical implementation of a SAR process is on the order of twenty to
5 thirty minutes, and thus the amount of data that may be lost when an RDF link and/or primary device fails could be twenty to thirty minutes worth of data.

Thus, it would be desirable to have an RDF system that exhibits some of the beneficial qualities of each of the different techniques discussed above while reducing the drawbacks. Such a system would exhibit low latency for each host write regardless of the
10 distance between the primary device and the secondary device and would provide consistency (recoverability) of the secondary device in case of failure.

Summary of the Invention

According to the present invention, recovering data provided in chunks to a plurality of secondary storage devices includes, for each of the secondary storage devices,
15 discarding data corresponding chunks for which all data thereof has not been received, and, for each of the secondary storage devices, restoring a chunk of data thereto where all of the chunks of data restored to the plurality of secondary storage devices correspond to a particular transmission cycle of primary storage devices that provide data to the plurality of secondary storage devices. Recovering data may also include, following
20 discarding and prior to restoring, for each of the plurality of secondary storage devices having two different chunks, waiting for external intervention to indicate whether to

restore a particular one of the chunks. The external intervention may be provided by a host computer that is proximate to at least one of the secondary storage devices or may be provided by a host computer that is proximate to at least one of the primary storage computers. Recovering data may also include restoring most recent chunks for all of the

5 plurality of secondary storage devices in response to there being two different chunks associated with all of the plurality of secondary storage devices, where a first one of the two chunks corresponds to a first transmission cycle and where a second one of the two chunks corresponding to a different transmission cycle. Recovering data may also include discarding chunks that are not restored. Recovering data may also include, for

10 each of the secondary storage devices, restoring a chunk of data corresponding to a particular transmission cycle where all of the secondary storage devices contain a chunk of data corresponding to the particular transmission cycle. Recovering data may also include discarding chunks that are not restored. Each transmission cycle may be assigned a particular tag value that is provided with each chunk of data. The tag values may be

15 used to determine the particular cycle for each of the chunks of data.

According further to the present invention, computer software that recovers data provided in chunks to a plurality of secondary storage devices includes executable code that discards data corresponding chunks for which all data thereof has not been received for each of the secondary storage devices and executable code that restores a chunk of

20 data thereto for each of the secondary storage devices, where all of the chunks of data restored to the plurality of secondary storage devices correspond to a particular transmission cycle of primary storage devices that provide data to the plurality of

secondary storage devices. The computer software may also include executable code that waits for external intervention to indicate whether to restore a particular one of the chunks for each of the plurality of secondary storage devices having two different chunks. The computer software may also include executable code that restores most

5 recent chunks for all of the plurality of secondary storage devices in response to there being two different chunks associated with all of the plurality of secondary storage devices, where a first one of the two chunks corresponds to a first transmission cycle and where a second one of the two chunks corresponding to a different transmission cycle.

The computer software may also include executable code that discards chunks that are

10 not restored. The computer software may also include executable code that restores a chunk of data corresponding to a particular transmission cycle where all of the secondary storage devices contain a chunk of data corresponding to the particular transmission cycle. The computer software may also include executable code that discards chunks that are not restored. Each transmission cycle may be assigned a particular tag value that is

15 provided with each chunk of data. The tag values may be used to determine the particular cycle for each of the chunks of data.

Brief Description of Drawings

Figure 1 is a schematic diagram showing a host, a local storage device, and a remote data storage device used in connection with the system described herein.

5 Figure 2 is a schematic diagram showing a flow of data between a host, a local storage device, and a remote data storage device used in connection with the system described herein.

Figure 3 is a schematic diagram illustrating items for constructing and manipulating chunks of data on a local storage device according to the system described herein.

10 Figure 4 is a diagram illustrating a data structure for a slot used in connection with the system described herein.

Figure 5 is a flow chart illustrating operation of a host adaptor (HA) in response to a write by a host according to the system described herein.

15 Figure 6 is a flow chart illustrating transferring data from a local storage device to a remote storage device according to the system described herein.

Figure 7 is a schematic diagram illustrating items for constructing and manipulating chunks of data on a remote storage device according to the system described herein.

Figure 8 is a flow chart illustrating steps performed by a remote storage device in connection with receiving a commit indicator from a local storage device according to the system described herein.
5

Figure 9 is a flow chart illustrating storing transmitted data at a remote storage device according to the system described herein.

Figure 10 is a flow chart illustrating steps performed in connection with a local storage device incrementing a sequence number according to a system described herein.
10

Figure 11 is a schematic diagram illustrating items for constructing and manipulating chunks of data on a local storage device according to an alternative embodiment of the system described herein.

Figure 12 is a flow chart illustrating operation of a host adaptor (HA) in response to a write by a host according to an alternative embodiment of the system described herein.
15

Figure 13 is a flow chart illustrating transferring data from a local storage device to a remote storage device according to an alternative embodiment of the system described herein.

5 Figure 14 is a schematic diagram illustrating a plurality of local and remote storage devices with a host according to the system described herein.

Figure 15 is a diagram showing a multi-box mode table used in connection with the system described herein.

Figure 16 is a flow chart illustrating modifying a multi-box mode table according to the system described herein.

10 Figure 17 is a flow chart illustrating cycle switching by the host according to the system described herein.

Figure 18 is a flow chart illustrating steps performed in connection with a local storage device incrementing a sequence number according to a system described herein.

15 Figure 19 is a flow chart illustrating transferring data from a local storage device to a remote storage device according to the system described herein.

Figure 20 is a flow chart illustrating transferring data from a local storage device to a remote storage device according to an alternative embodiment of the system described herein.

Figure 21 is a flow chart illustrating providing an active empty indicator message
5 from a remote storage device to a corresponding local storage device according to the system described herein.

Figure 22 is a schematic diagram illustrating a plurality of local and remote storage devices with a plurality of hosts according to the system described herein.

Figure 23 is a flow chart illustrating a processing performed by a remote storage
10 device in connection with data recovery according to the system described herein.

Figure 24 is a flow chart illustrating a processing performed by a host in connection with data recovery according to the system described herein.

Detailed Description of Various Embodiments

Referring to Figure 1, a diagram 20 shows a relationship between a host 22, a
15 local storage device 24 and a remote storage device 26. The host 22 reads and writes data from and to the local storage device 24 via a host adapter (HA) 28, which facilitates the interface between the host 22 and the local storage device 24. Although the diagram 20 only shows one host 22 and one HA 28, it will be appreciated by one of ordinary skill in

the art that multiple HA's may be used and that one or more HA's may have one or more hosts coupled thereto.

Data from the local storage device 24 is copied to the remote storage device 26 via an RDF link 29 to cause the data on the remote storage device 26 to be identical to the 5 data on the local storage device 24. Although only the one link 29 is shown, it is possible to have additional links between the storage devices 24, 26 and to have links between one or both of the storage devices 24, 26 and other storage devices (not shown). Note that there may be a time delay between the transfer of data from the local storage device 24 to the remote storage device 26, so that the remote storage device 26 may, at certain points 10 in time, contain data that is not identical to the data on the local storage device 24.

Communication using RDF is described, for example, in U.S. Patent No. 5,742,792, which is incorporated by reference herein.

The local storage device 24 includes a first plurality of RDF adapter units (RA's) 30a, 30b, 30c and the remote storage device 26 includes a second plurality of RA's 32a- 15 32c. The RA's 30a-30c, 32a-32c are coupled to the RDF link 29 and are similar to the host adapter 28, but are used to transfer data between the storage devices 24, 26. The software used in connection with the RA's 30a-30c, 32a-32c is discussed in more detail hereinafter.

The storage devices 24, 26 may include one or more disks, each containing a 20 different portion of data stored on each of the storage devices 24, 26. Figure 1 shows the

storage device 24 including a plurality of disks 33a, 33b, 33c and the storage device 26 including a plurality of disks 34a, 34b, 34c. The RDF functionality described herein may be applied so that the data for at least a portion of the disks 33a-33c of the local storage device 24 is copied, using RDF, to at least a portion of the disks 34a-34c of the remote storage device 26. It is possible that other data of the storage devices 24, 26 is not copied between the storage devices 24, 26, and thus is not identical.

5 Each of the disks 33a-33c is coupled to a corresponding disk adapter unit (DA) 35a, 35b, 35c that provides data to a corresponding one of the disks 33a-33c and receives data from a corresponding one of the disks 33a-33c. Similarly, a plurality of DA's 36a, 10 36b, 36c of the remote storage device 26 are used to provide data to corresponding ones of the disks 34a-34c and receive data from corresponding ones of the disks 34a-34c. An internal data path exists between the DA's 35a-35c, the HA 28 and the RA's 30a-30c of the local storage device 24. Similarly, an internal data path exists between the DA's 36a- 15 36c and the RA's 32a-32c of the remote storage device 26. Note that, in other embodiments, it is possible for more than one disk to be serviced by a DA and that it is possible for more than one DA to service a disk.

The local storage device 24 also includes a global memory 37 that may be used to facilitate data transferred between the DA's 35a-35c, the HA 28 and the RA's 30a-30c. The memory 37 may contain tasks that are to be performed by one or more of the DA's 20 35a-35c, the HA 28 and the RA's 30a-30c, and a cache for data fetched from one or more of the disks 33a-33c. Similarly, the remote storage device 26 includes a global memory

38 that may contain tasks that are to be performed by one or more of the DA's 36a-36c and the RA's 32a-32c, and a cache for data fetched from one or more of the disks 34a-34c. Use of the memories 37, 38 is described in more detail hereinafter.

The storage space in the local storage device 24 that corresponds to the disks 5 33a-33c may be subdivided into a plurality of volumes or logical devices. The logical devices may or may not correspond to the physical storage space of the disks 33a-33c. Thus, for example, the disk 33a may contain a plurality of logical devices or, alternatively, a single logical device could span both of the disks 33a, 33b. Similarly, the storage space for the remote storage device 26 that comprises the disks 34a-34c may be 10 subdivided into a plurality of volumes or logical devices, where each of the logical devices may or may not correspond to one or more of the disks 34a-34c.

Providing an RDF mapping between portions of the local storage device 24 and the remote storage device 26 involves setting up a logical device on the remote storage device 26 that is a remote mirror for a logical device on the local storage device 24. The 15 host 22 reads and writes data from and to the logical device on the local storage device 24 and the RDF mapping causes modified data to be transferred from the local storage device 24 to the remote storage device 26 using the RA's, 30a-30c, 32a-32c and the RDF link 29. In steady state operation, the logical device on the remote storage device 26 contains data that is identical to the data of the logical device on the local storage device 20 24. The logical device on the local storage device 24 that is accessed by the host 22 is referred to as the "R1 volume" (or just "R1") while the logical device on the remote

storage device 26 that contains a copy of the data on the R1 volume is called the "R2 volume" (or just "R2"). Thus, the host reads and writes data from and to the R1 volume and RDF handles automatic copying and updating of the data from the R1 volume to the R2 volume.

5 Referring to Figure 2, a path of data is illustrated from the host 22 to the local storage device 24 and the remote storage device 26. Data written from the host 22 to the local storage device 24 is stored locally, as illustrated by the data element 51 of the local storage device 24. The data that is written by the host 22 to the local storage device 24 is also maintained by the local storage device 24 in connection with being sent by the local 10 storage device 24 to the remote storage device 26 via the link 29.

In the system described herein, each data write by the host 22 (of, for example a record, a plurality of records, a track, etc.) is assigned a sequence number. The sequence number may be provided in an appropriate data field associated with the write. In Figure 2, the writes by the host 22 are shown as being assigned sequence number N. All of the 15 writes performed by the host 22 that are assigned sequence number N are collected in a single chunk of data 52. The chunk 52 represents a plurality of separate writes by the host 22 that occur at approximately the same time.

Generally, the local storage device 24 accumulates chunks of one sequence number while transmitting a previously accumulated chunk (having the previous 20 sequence number) to the remote storage device 26. Thus, while the local storage device

24 is accumulating writes from the host 22 that are assigned sequence number N, the
writes that occurred for the previous sequence number (N-1) are transmitted by the local
storage device 24 to the remote storage device 26 via the link 29. A chunk 54 represents
writes from the host 22 that were assigned the sequence number N-1 that have not been
5 transmitted yet to the remote storage device 26.

The remote storage device 26 receives the data from the chunk 54 corresponding
to writes assigned a sequence number N-1 and constructs a new chunk 56 of host writes
having sequence number N-1. The data may be transmitted using appropriate RDF
protocol that acknowledges data sent across the link 29. When the remote storage device
10 26 has received all of the data from the chunk 54, the local storage device 24 sends a
commit message to the remote storage device 26 to commit all the data assigned the N-1
sequence number corresponding to the chunk 56. Generally, once a chunk corresponding
to a particular sequence number is committed, that chunk may be written to the logical
storage device. This is illustrated in Figure 2 with a chunk 58 corresponding to writes
15 assigned sequence number N-2 (i.e., two before the current sequence number being used
in connection with writes by the host 22 to the local storage device 26). In Figure 2, the
chunk 58 is shown as being written to a data element 62 representing disk storage for the
remote storage device 26. Thus, the remote storage device 26 is receiving and
accumulating the chunk 56 corresponding to sequence number N-1 while the chunk 58
20 corresponding to the previous sequence number (N-2) is being written to disk storage of
the remote storage device 26 illustrated by the data element 62. In some embodiments,

the data for the chunk 58 is marked for write (but not necessarily written immediately), while the data for the chunk 56 is not.

Thus, in operation, the host 22 writes data to the local storage device 24 that is stored locally in the data element 51 and is accumulated in the chunk 52. Once all of the 5 data for a particular sequence number has been accumulated (described elsewhere herein), the local storage device 24 increments the sequence number. Data from the chunk 54 corresponding to one less than the current sequence number is transferred from the local storage device 24 to the remote storage device 26 via the link 29. The chunk 58 corresponds to data for a sequence number that was committed by the local storage 10 device 24 sending a message to the remote storage device 26. Data from the chunk 58 is written to disk storage of the remote storage device 26.

Note that the writes within a particular one of the chunks 52, 54, 56, 58 are not necessarily ordered. However, as described in more detail elsewhere herein, every write for the chunk 58 corresponding to sequence number N-2 was begun prior to beginning 15 any of the writes for the chunks 54, 56 corresponding to sequence number N-1. In addition, every write for the chunks 54, 56 corresponding to sequence number N-1 was begun prior to beginning any of the writes for the chunk 52 corresponding to sequence number N. Thus, in the event of a communication failure between the local storage device 24 and the remote storage device 26, the remote storage device 26 may simply 20 finish writing the last committed chunk of data (the chunk 58 in the example of Figure 2) and can be assured that the state of the data at the remote storage device 26 is ordered in

the sense that the data element 62 contains all of the writes that were begun prior to a certain point in time and contains no writes that were begun after that point in time.

Thus, R2 always contains a point in time copy of R1 and it is possible to reestablish a consistent image from the R2 device.

5 Referring to Figure 3, a diagram 70 illustrates items used to construct and maintain the chunks 52, 54. A standard logical device 72 contains data written by the host 22 and corresponds to the data element 51 of Figure 2 and the disks 33a-33c of Figure 1. The standard logical device 72 contains data written by the host 22 to the local storage device 24.

10 Two linked lists of pointers 74, 76 are used in connection with the standard logical device 72. The linked lists 74, 76 correspond to data that may be stored, for example, in the memory 37 of the local storage device 24. The linked list 74 contains a plurality of pointers 81-85, each of which points to a slot of a cache 88 used in connection with the local storage device 24. Similarly, the linked list 76 contains a
15 plurality of pointers 91-95, each of which points to a slot of the cache 88. In some embodiments, the cache 88 may be provided in the memory 37 of the local storage device 24. The cache 88 contains a plurality of cache slots 102-104 that may be used in connection to writes to the standard logical device 72 and, at the same time, used in connection with the linked lists 74, 76.

Each of the linked lists 74, 76 may be used for one of the chunks of data 52, 54 so that, for example, the linked list 74 may correspond to the chunk of data 52 for sequence number N while the linked list 76 may correspond to the chunk of data 54 for sequence number N-1. Thus, when data is written by the host 22 to the local storage device 24, the
5 data is provided to the cache 88 and, in some cases (described elsewhere herein), an appropriate pointer of the linked list 74 is created. Note that the data will not be removed from the cache 88 until the data is destaged to the standard logical device 72 and the data is also no longer pointed to by one of the pointers 81-85 of the linked list 74, as described elsewhere herein.

10 In an embodiment herein, one of the linked lists 74, 76 is deemed “active” while the other is deemed “inactive”. Thus, for example, when the sequence number N is even, the linked list 74 may be active while the linked list 76 is inactive. The active one of the linked lists 74, 76 handles writes from the host 22 while the inactive one of the linked lists 74, 76 corresponds to the data that is being transmitted from the local storage device
15 24 to the remote storage device 26.

While the data that is written by the host 22 is accumulated using the active one of the linked lists 74, 76 (for the sequence number N), the data corresponding to the inactive one of the linked lists 74, 76 (for previous sequence number N-1) is transmitted from the local storage device 24 to the remote storage device 26. The RA’s 30a-30c use the linked
20 lists 74, 76 to determine the data to transmit from the local storage device 24 to the remote storage device 26.

Once data corresponding to a particular one of the pointers in one of the linked lists 74, 76 has been transmitted to the remote storage device 26, the particular one of the pointers may be removed from the appropriate one of the linked lists 74, 76. In addition, the data may also be marked for removal from the cache 88 (i.e., the slot may be returned 5 to a pool of slots for later, unrelated, use) provided that the data in the slot is not otherwise needed for another purpose (e.g., to be destaged to the standard logical device 72). A mechanism may be used to ensure that data is not removed from the cache 88 until all devices are no longer using the data. Such a mechanism is described, for example, in U.S. Patent No. 5,537,568 issued on July 16, 1996 and in U.S. patent 10 application no. 09/850,551 filed on July 7, 2001, both of which are incorporated by reference herein.

Referring to Figure 4, a slot 120, like one of the slots 102-104 of the cache 88, includes a header 122 and data 124. The header 122 corresponds to overhead information used by the system to manage the slot 120. The data 124 is the 15 corresponding data from the disk that is being (temporarily) stored in the slot 120. Information in the header 122 includes pointers back to the disk, time stamp(s), etc.

The header 122 also includes a cache stamp 126 used in connection with the system described herein. In an embodiment herein, the cache stamp 126 is eight bytes. Two of the bytes are a “password” that indicates whether the slot 120 is being used by the 20 system described herein. In other embodiments, the password may be one byte while the following byte is used for a pad. As described elsewhere herein, the two bytes of the

password (or one byte, as the case may be) being equal to a particular value indicates that the slot 120 is pointed to by at least one entry of the linked lists 74, 76. The password not being equal to the particular value indicates that the slot 120 is not pointed to by an entry of the linked lists 74, 76. Use of the password is described elsewhere herein.

- 5 The cache stamp 126 also includes a two byte field indicating the sequence number (e.g., N, N-1, N-2, etc.) of the data 124 of the slot 120. As described elsewhere herein, the sequence number field of the cache stamp 126 may be used to facilitate the processing described herein. The remaining four bytes of the cache stamp 126 may be used for a pointer, as described elsewhere herein. Of course, the two bytes of the
- 10 sequence number and the four bytes of the pointer are only valid when the password equals the particular value that indicates that the slot 120 is pointed to by at least one entry in one of the lists 74, 76.

- Referring to Figure 5, a flow chart 140 illustrates steps performed by the HA 28 in connection with a host 22 performing a write operation. Of course, when the host 22
- 15 performs a write, processing occurs for handling the write in a normal fashion irrespective of whether the data is part of an R1/R2 RDF group. For example, when the host 22 writes data for a portion of the disk, the write occurs to a cache slot which is eventually destaged to the disk. The cache slot may either be a new cache slot or may be an already existing cache slot created in connection with a previous read and/or write
- 20 operation to the same track.

Processing begins at a first step 142 where a slot corresponding to the write is locked. In an embodiment herein, each of the slots 102-104 of the cache 88 corresponds to a track of data on the standard logical device 72. Locking the slot at the step 142 prevents additional processes from operating on the relevant slot during the processing 5 performed by the HA 28 corresponding to the steps of the flow chart 140.

Following step 142 is a step 144 where a value for N, the sequence number, is set. As discussed elsewhere herein, the value for the sequence number obtained at the step 144 is maintained during the entire write operation performed by the HA 28 while the slot is locked. As discussed elsewhere herein, the sequence number is assigned to each 10 write to set the one of the chunks of data 52, 54 to which the write belongs. Writes performed by the host 22 are assigned the current sequence number. It is useful that a single write operation maintain the same sequence number throughout.

Following the step 144 is a test step 146 which determines if the password field of the cache slot is valid. As discussed above, the system described herein sets the 15 password field to a predetermined value to indicate that the cache slot is already in one of the linked lists of pointers 74, 76. If it is determined at the test step 146 that the password field is not valid (indicating that the slot is new and that no pointers from the lists 74, 76 point to the slot), then control passes from the step 146 to a step 148, where the cache stamp of the new slot is set by setting the password to the predetermined value, setting 20 the sequence number field to N, and setting the pointer field to Null. In other embodiments, the pointer field may be set to point to the slot itself.

Following the step 148 is a step 152 where a pointer to the new slot is added to the active one of the pointer lists 74, 76. In an embodiment herein, the lists 74, 76 are circular doubly linked lists, and the new pointer is added to the circular doubly linked list in a conventional fashion. Of course, other appropriate data structures could be used to

5 manage the lists 74, 76. Following the step 152 is a step 154 where flags are set. At the step 154, the RDF_WP flag (RDF write pending flag) is set to indicate that the slot needs to be transmitted to the remote storage device 26 using RDF. In addition, at the step 154, the IN_CACHE flag is set to indicate that the slot needs to be destaged to the standard logical device 72. Following the step 154 is a step 156 where the data being written by

10 the host 22 and the HA 28 is written to the slot. Following the step 156 is a step 158 where the slot is unlocked. Following step 158, processing is complete.

If it is determined at the test step 146 that the password field of the slot is valid (indicating that the slot is already pointed to by at least one pointer of the lists 74, 76), then control transfers from the step 146 to a test step 162, where it is determined whether

15 the sequence number field of the slot is equal to the current sequence number, N. Note that there are two valid possibilities for the sequence number field of a slot with a valid password. It is possible for the sequence number field to be equal to N, the current sequence number. This occurs when the slot corresponds to a previous write with sequence number N. The other possibility is for the sequence number field to equal N-1.

20 This occurs when the slot corresponds to a previous write with sequence number N-1. Any other value for the sequence number field is invalid. Thus, for some embodiments, it may be possible to include error/validity checking in the step 162 or possibly make

error/validity checking a separate step. Such an error may be handled in any appropriate fashion, which may include providing a message to a user.

If it is determined at the step 162 that the value in the sequence number field of the slot equals the current sequence number N, then no special processing is required and

5 control transfers from the step 162 to the step 156, discussed above, where the data is written to the slot. Otherwise, if the value of the sequence number field is N-1 (the only other valid value), then control transfers from the step 162 to a step 164 where a new slot is obtained. The new slot obtained at the step 164 may be used to store the data being written.

10 Following the step 164 is a step 166 where the data from the old slot is copied to the new slot that was obtained at the step 164. Note that that the copied data includes the RDF_WP flag, which should have been set at the step 154 on a previous write when the slot was first created. Following the step 166 is a step 168 where the cache stamp for the new slot is set by setting the password field to the appropriate value, setting the sequence

15 number field to the current sequence number, N, and setting the pointer field to point to the old slot. Following the step 168 is a step 172 where a pointer to the new slot is added to the active one of the linked lists 74, 76. Following the step 172 is the step 156, discussed above, where the data is written to the slot which, in this case, is the new slot.

20 Referring to Figure 6, a flow chart 200 illustrates steps performed in connection with the RA's 30a-30c scanning the inactive one of the lists 72, 74 to transmit RDF data

from the local storage device 24 to the remote storage device 26. As discussed above, the inactive one of the lists 72, 74 points to slots corresponding to the N-1 cycle for the R1 device when the N cycle is being written to the R1 device by the host using the active one of the lists 72, 74.

5 Processing begins at a first step 202 where it is determined if there are any entries in the inactive one of the lists 72, 74. As data is transmitted, the corresponding entries are removed from the inactive one of the lists 72, 74. In addition, new writes are provided to the active one of the lists 72, 74 and not generally to the inactive one of the lists 72, 74. Thus, it is possible (and desirable, as described elsewhere herein) for the
10 inactive one of the lists 72, 74 to contain no data at certain times. If it is determined at the step 202 that there is no data to be transmitted, then the inactive one of the lists 72, 74 is continuously polled until data becomes available. Data for sending becomes available in connection with a cycle switch (discussed elsewhere herein) where the inactive one of the lists 72, 74 becomes the active one of the lists 72, 74, and vice versa.

15 If it is determined at the step 202 that there is data available for sending, control transfers from the step 202 to a step 204, where the slot is verified as being correct. The processing performed at the step 204 is an optional "sanity check" that may include verifying that the password field is correct and verifying that the sequence number field is correct. If there is incorrect (unexpected) data in the slot, error processing may be
20 performed, which may include notifying a user of the error and possibly error recovery processing.

Following the step 204 is a step 212, where the data is sent via RDF in a conventional fashion. In an embodiment herein, the entire slot is not transmitted. Rather, only records within the slot that have the appropriate mirror bits set (indicating the records have changed) are transmitted to the remote storage device 26. However, in other 5 embodiments, it may be possible to transmit the entire slot, provided that the remote storage device 26 only writes data corresponding to records having appropriate mirror bits set and ignores other data for the track, which may or may not be valid. Following the step 212 is a test step 214 where it is determined if the data that was transmitted has been acknowledged by the R2 device. If not, the data is resent, as indicated by the flow 10 from the step 214 back to the step 212. In other embodiments, different and more involved processing may used to send data and acknowledge receipt thereof. Such processing may include error reporting and alternative processing that is performed after a certain number of attempts to send the data have failed.

Once it is determined at the test step 214 that the data has been successfully sent, 15 control passes from the step 214 to a step 216 to clear the RDF_WP flag (since the data has been successfully sent via RDF). Following the step 216 is a test step 218 where it is determined if the slot is a duplicate slot created in connection with a write to a slot already having an existing entry in the inactive one of the lists 72, 74. This possibility is discussed above in connection with the steps 162, 164, 166, 168, 172. If it is determined 20 at the step 218 that the slot is a duplicate slot, then control passes from the step 218 to a step 222 where the slot is returned to the pool of available slots (to be reused). In addition, the slot may also be aged (or have some other appropriate mechanism applied

thereto) to provide for immediate reuse ahead of other slots since the data provided in the slot is not valid for any other purpose. Following the step 222 or the step 218 if the slot is not a duplicate slot is a step 224 where the password field of the slot header is cleared so that when the slot is reused, the test at the step 146 of Figure 5 properly classifies the
5 slot as a new slot.

Following the step 224 is a step 226 where the entry in the inactive one of the lists 72, 74 is removed. Following the step 226, control transfers back to the step 202, discussed above, where it is determined if there are additional entries on the inactive one of the lists 72, 74 corresponding to data needing to be transferred.

10 Referring to Figure 7, a diagram 240 illustrates creation and manipulation of the chunks 56, 58 used by the remote storage device 26. Data that is received by the remote storage device 26, via the link 29, is provided to a cache 242 of the remote storage device 26. The cache 242 may be provided, for example, in the memory 38 of the remote storage device 26. The cache 242 includes a plurality of cache slots 244-246, each of
15 which may be mapped to a track of a standard logical storage device 252. The cache 242 is similar to the cache 88 of Figure 3 and may contain data that can be destaged to the standard logical storage device 252 of the remote storage device 26. The standard logical storage device 252 corresponds to the data element 62 shown in Figure 2 and the disks 34a-34c shown in Figure 1.

The remote storage device 26 also contains a pair of cache only virtual devices 254, 256. The cache only virtual devices 254, 256 corresponded device tables that may be stored, for example, in the memory 38 of the remote storage device 26. Each track entry of the tables of each of the cache only virtual devices 254, 256 point to either a 5 track of the standard logical device 252 or point to a slot of the cache 242. Cache only virtual devices are described in a copending U.S. patent application titled CACHE-ONLY VIRTUAL DEVICES, filed on March 25, 2003 and having serial no. 10/396,800, which is incorporated by reference herein.

The plurality of cache slots 244-246 may be used in connection to writes to the 10 standard logical device 252 and, at the same time, used in connection with the cache only virtual devices 254, 256. In an embodiment herein, each of track table entry of the cache only virtual devices 254, 256 contain a null to indicate that the data for that track is stored on a corresponding track of the standard logical device 252. Otherwise, an entry in the track table for each of the cache only virtual devices 254, 256 contains a pointer to one of 15 the slots 244-246 in the cache 242.

Each of the cache only virtual devices 254, 256 corresponds to one of the data chunks 56, 58. Thus, for example, the cache only virtual device 254 may correspond to the data chunk 56 while the cache only virtual device 256 may correspond to the data chunk 58. In an embodiment herein, one of the cache only virtual devices 254, 256 may 20 be deemed “active” while the other one of the cache only virtual devices 254, 256 may be deemed “inactive”. The inactive one of the cache only virtual devices 254, 256 may

correspond to data being received from the local storage device 24 (i.e., the chunk 56) while the active one of the cache only virtual device 254, 256 corresponds to data being restored (written) to the standard logical device 252.

Data from the local storage device 24 that is received via the link 29 may be
5 placed in one of the slots 244-246 of the cache 242. A corresponding pointer of the inactive one of the cache only virtual devices 254, 256 may be set to point to the received data. Subsequent data having the same sequence number may be processed in a similar manner. At some point, the local storage device 24 provides a message committing all of the data sent using the same sequence number. Once the data for a particular sequence
10 number has been committed, the inactive one of the cache only virtual devices 254, 256 becomes active and vice versa. At that point, data from the now active one of the cache only virtual devices 254, 256 is copied to the standard logical device 252 while the inactive one of the cache only virtual devices 254, 256 is used to receive new data (having a new sequence number) transmitted from the local storage device 24 to the
15 remote storage device 26.

As data is removed from the active one of the cache only virtual devices 254, 256 (discussed elsewhere herein), the corresponding entry in the active one of the cache only virtual devices 254, 256 may be set to null. In addition, the data may also be removed from the cache 244 (i.e., the slot returned to the pool of free slots for later use) provided
20 that the data in the slot is not otherwise needed for another purpose (e.g., to be destaged to the standard logical device 252). A mechanism may be used to ensure that data is not

removed from the cache 242 until all mirrors (including the cache only virtual devices 254, 256) are no longer using the data. Such a mechanism is described, for example, in U.S. Patent No. 5,537,568 issued on July 16, 1996 and in U.S. patent application no. 09/850,551 filed on July 7, 2001, both of which are incorporated by reference herein.

5 In some embodiments discussed elsewhere herein, the remote storage device 26 may maintain linked lists 258, 262 like the lists 74, 76 used by the local storage device 24. The lists 258, 262 may contain information that identifies the slots of the corresponding cache only virtual devices 254, 256 that have been modified, where one of the lists 258, 262 corresponds to one of the cache only virtual devices 254, 256 and the other one of the lists 258, 262 corresponds to the other one of the cache only virtual devices 254, 256. As discussed elsewhere herein, the lists 258, 262 may be used to facilitate restoring data from the cache only virtual devices 254, 256 to the standard logical device 252.

10

15 Referring to Figure 8, a flow chart 270 illustrates steps performed by the remote storage device 26 in connection with processing data for a sequence number commit transmitted by the local storage device 24 to the remote storage device 26. As discussed elsewhere herein, the local storage device 24 periodically increments sequence numbers. When this occurs, the local storage device 24 finishes transmitting all of the data for the previous sequence number and then sends a commit message for the previous sequence
20 number.

Processing begins at a first step 272 where the commit is received. Following the step 272 is a test step 274 which determines if the active one of the cache only virtual devices 254, 256 of the remote storage device 26 is empty. As discussed elsewhere herein, the inactive one of the cache only virtual devices 254, 256 of the remote storage 5 device 26 is used to accumulate data from the local storage device 24 sent using RDF while the active one of the cache only virtual devices 254, 256 is restored to the standard logical device 252.

- If it is determined at the test step 274 that the active one of the cache only virtual devices 254, 256 is not empty, then control transfers from the test step 274 to a step 276 10 where the restore for the active one of the cache only virtual devices 254, 256 is completed prior to further processing being performed. Restoring data from the active one of the cache only virtual devices 254, 256 is described in more detail elsewhere herein. It is useful that the active one of the cache only virtual devices 254, 256 is empty prior to handling the commit and beginning to restore data for the next sequence number.
- 15 Following the step 276 or following the step 274 if the active one of the cache only virtual devices 254, 256 is determined to be empty, is a step 278 where the active one of the cache only virtual devices 254, 256 is made inactive. Following the step 278 is a step 282 where the previously inactive one of the cache only virtual devices 254, 256 (i.e., the one that was inactive prior to execution of the step 278) is made active.
- 20 Swapping the active and inactive cache only virtual devices 254, 256 at the steps 278,

282 prepares the now inactive (and empty) one of the cache only virtual devices 254, 256 to begin to receive data from the local storage device 24 for the next sequence number.

Following the step 282 is a step 284 where the active one of the cache only virtual devices 254, 256 is restored to the standard logical device 252 of the remote storage

- 5 device 26. Restoring the active one of the cache only virtual devices 254, 256 to the standard logical device 252 is described in more detail hereinafter. However, note that, in some embodiments, the restore process is begun, but not necessarily completed, at the step 284. Following the step 284 is a step 286 where the commit that was sent from the local storage device 24 to the remote storage device 26 is acknowledged back to the local
10 storage device 24 so that the local storage device 24 is informed that the commit was successful. Following the step 286, processing is complete.

Referring to Figure 9, a flow chart 300 illustrates in more detail the steps 276, 284 of Figure 8 where the remote storage device 26 restores the active one of the cache only virtual devices 254, 256. Processing begins at a first step 302 where a pointer is set to

- 15 point to the first slot of the active one of the cache only virtual devices 254, 256. The pointer is used to iterate through each track table entry of the active one of the cache only virtual devices 254, 256, each of which is processed individually. Following the step 302 is a test step 304 where it is determined if the track of the active one of the cache only virtual devices 254, 256 that is being processed points to the standard logical device 252.
20 If so, then there is nothing to restore. Otherwise, control transfers from the step 304 to a

step a 306 where the corresponding slot of the active one of the cache only virtual devices 254, 256 is locked.

Following the step 306 is a test step 308 which determines if the corresponding slot of the standard logical device 252 is already in the cache of the remote storage device

- 5 26. If so, then control transfers from the test step 308 to a step 312 where the slot of the standard logical device is locked. Following step 312 is a step 314 where the data from the active one of the cache only virtual devices 254, 256 is merged with the data in the cache for the standard logical device 252. Merging the data at the step 314 involves overwriting the data for the standard logical device with the new data of the active one of
- 10 the cache only virtual devices 254, 256. Note that, in embodiments that provide for record level flags, it may be possible to simply OR the new records from the active one of the cache only virtual devices 254, 256 to the records of the standard logical device 252 in the cache. That is, if the records are interleaved, then it is only necessary to use the records from the active one of the cache only virtual devices 254, 256 that have changed
- 15 and provide the records to the cache slot of the standard logical device 252. Following step 314 is a step 316 where the slot of the standard logical device 252 is unlocked. Following step 316 is a step 318 where the slot of the active one of the cache only virtual devices 254, 256 that is being processed is also unlocked.

If it is determined at the test step 308 that the corresponding slot of the standard

- 20 logical device 252 is not in cache, then control transfers from the test step 308 to a step 322 where the track entry for the slot of the standard logical device 252 is changed to

indicate that the slot of the standard logical device 252 is in cache (e.g., an IN_CACHE flag may be set) and needs to be destaged. As discussed elsewhere herein, in some embodiments, only records of the track having appropriate mirror bits set may need to be destaged. Following the step 322 is a step 324 where a flag for the track may be set to indicate that the data for the track is in the cache.

Following the step 324 is a step 326 where the slot pointer for the standard logical device 252 is changed to point to the slot in the cache. Following the step 326 is a test step 328 which determines if the operations performed at the steps 322, 324, 326 have been successful. In some instances, a single operation called a "compare and swap" operation may be used to perform the steps 322, 324, 326. If these operations are not successful for any reason, then control transfers from the step 328 back to the step 308 to reexamine if the corresponding track of the standard logical device 252 is in the cache. Otherwise, if it is determined at the test step 328 that the previous operations have been successful, then control transfers from the test step 328 to the step 318, discussed above.

Following the step 318 is a test step 332 which determines if the cache slot of the active one of the cache only virtual devices 254, 256 (which is being restored) is still being used. In some cases, it is possible that the slot for the active one of the cache only virtual devices 254, 256 is still being used by another mirror. If it is determined at the test step 332 that the slot of the cache only virtual device is not being used by another mirror, then control transfers from the test step 332 to a step 334 where the slot is released for use by other processes (e.g., restored to pool of available slots, as discussed

elsewhere herein). Following the step 334 is a step 336 to point to the next slot to process the next slot of the active one of the cache only virtual devices 254, 256. Note that the step 336 is also reached from the test step 332 if it is determined at the step 332 that the active one of the cache only virtual devices 254, 256 is still being used by another mirror. Note also that the step 336 is reached from the test step 304 if it is determined at the step 304 that, for the slot being processed, the active one of the cache only virtual devices 254, 256 points to the standard logical device 252. Following the step 336 is a test step 338 which determines if there are more slots of the active one of the cache only virtual devices 254, 256 to be processed. If not, processing is complete. Otherwise,

5 control transfers from the test step 338 back to the step 304.

10

In another embodiment, it is possible to construct lists of modified slots for the received chunk of data 56 corresponding to the N-1 cycle on the remote storage device 26, such as the lists 258, 262 shown in Figure 7. As the data is received, the remote storage device 26 constructs a linked list of modified slots. The lists that are constructed

15 may be circular, linear (with a NULL termination), or any other appropriate design. The lists may then be used to restore the active one of the cache only virtual devices 254, 256.

The flow chart 300 of Figure 9 shows two alternative paths 342, 344 that illustrate operation of embodiments where a list of modified slots is used. At the step 302, a pointer (used for iterating through the list of modified slots) is made to point to the first

20 element of the list. Following the step 302 is the step 306, which is reached by the

alternative path 342. In embodiments that use lists of modified slots, the test step 304 is not needed since no slots on the list should point to the standard logical device 252.

Following the step 306, processing continues as discussed above with the previous embodiment, except that the step 336 refers to traversing the list of modified 5 slots rather than pointing to the next slot in the COVD. Similarly, the test at the step 338 determines if the pointer is at the end of the list (or back to the beginning in the case of a circular linked list). Also, if it is determined at the step 338 that there are more slots to process, then control transfers from the step 338 to the step 306, as illustrated by the alternative path 344. As discussed above, for embodiments that use a list of modified 10 slots, the step 304 may be eliminated.

Referring to Figure 10, a flow chart 350 illustrates steps performed in connection with the local storage device 24 increasing the sequence number. Processing begins at a first step 352 where the local storage device 24 waits at least M seconds prior to increasing the sequence number. In an embodiment herein, M is thirty, but of course M 15 could be any number. Larger values for M increase the amount of data that may be lost if communication between the storage devices 24, 26 is disrupted. However, smaller values for M increase the total amount of overhead caused by incrementing the sequence number more frequently.

Following the step 352 is a test step 354 which determines if all of the HA's of 20 the local storage device 24 have set a bit indicating that the HA's have completed all of

the I/O's for a previous sequence number. When the sequence number changes, each of the HA's notices the change and sets a bit indicating that all I/O's of the previous sequence number are completed. For example, if the sequence number changes from N-1 to N, an HA will set the bit when the HA has completed all I/O's for sequence number N-

- 5 1. Note that, in some instances, a single I/O for an HA may take a long time and may still be in progress even after the sequence number has changed. Note also that, for some systems, a different mechanism may be used to determine if all of the HA's have completed their N-1 I/O's. The different mechanism may include examining device tables in the memory 37.

- 10 If it is determined at the test step 354 that I/O's from the previous sequence number have been completed, then control transfers from the step 354 to a test step 356 which determines if the inactive one of the lists 74, 76 is empty. Note that a sequence number switch may not be made unless and until all of the data corresponding to the inactive one of the lists 74, 76 has been completely transmitted from the local storage
15 device 24 to the remote storage device 26 using the RDF protocol. Once the inactive one of the lists 74, 76 is determined to be empty, then control transfers from the step 356 to a step 358 where the commit for the previous sequence number is sent from the local storage device 24 to the remote storage device 26. As discussed above, the remote storage device 26 receiving a commit message for a particular sequence number will
20 cause the remote storage device 26 to begin restoring the data corresponding to the sequence number.

Following the step 358 is a step 362 where the copying of data for the inactive one of the lists 74, 76 is suspended. As discussed elsewhere herein, the inactive one of the lists is scanned to send corresponding data from the local storage device 24 to the remote storage device 26. It is useful to suspend copying data until the sequence number 5 switch is completed. In an embodiment herein, the suspension is provided by sending a message to the RA's 30a-30c. However, it will be appreciated by one of ordinary skill in the art that for embodiments that use other components to facilitate sending data using the system described herein, suspending copying may be provided by sending appropriate messages/commands to the other components.

- 10 Following step 362 is a step 364 where the sequence number is incremented. Following step 364 is a step 366 where the bits for the HA's that are used in the test step 354 are all cleared so that the bits may be set again in connection with the increment of the sequence number. Following step 366 is a test step 372 which determines if the remote storage device 26 has acknowledged the commit message sent at the step 358.
- 15 Acknowledging the commit message is discussed above in connection with Figure 8. Once it is determined that the remote storage device 26 has acknowledged the commit message sent at the step 358, control transfers from the step 372 to a step 374 where the suspension of copying, which was provided at the step 362, is cleared so that copying may resume. Following step 374, processing is complete. Note that it is possible to go 20 from the step 374 back to the step 352 to begin a new cycle to continuously increment the sequence number.

It is also possible to use COVD's on the R1 device to collect slots associated with active data and inactive chunks of data. In that case, just as with the R2 device, one COVD could be associated with the inactive sequence number and another COVD could be associated with the active sequence number. This is described below.

5 Referring to Figure 11, a diagram 400 illustrates items used to construct and maintain the chunks 52, 54. A standard logical device 402 contains data written by the host 22 and corresponds to the data element 51 of Figure 2 and the disks 33a-33c of Figure 1. The standard logical device 402 contains data written by the host 22 to the local storage device 24.

10 Two cache only virtual devices 404, 406 are used in connection with the standard logical device 402. The cache only virtual devices 404, 406 corresponded device tables that may be stored, for example, in the memory 37 of the local storage device 24. Each track entry of the tables of each of the cache only virtual devices 404, 406 point to either a track of the standard logical device 402 or point to a slot of a cache 408 used in
15 connection with the local storage device 24. In some embodiments, the cache 408 may be provided in the memory 37 of the local storage device 24.

The cache 408 contains a plurality of cache slots 412-414 that may be used in connection to writes to the standard logical device 402 and, at the same time, used in connection with the cache only virtual devices 404, 406. In an embodiment herein, each
20 track table entry of the cache only virtual devices 404, 406 contains a null to point to a

corresponding track of the standard logical device 402. Otherwise, an entry in the track table for each of the cache only virtual devices 404, 406 contains a pointer to one of the slots 412-414 in the cache 408.

Each of the cache only virtual devices 404, 406 may be used for one of the chunks 5 of data 52, 54 so that, for example, the cache only virtual device 404 may correspond to the chunk of data 52 for sequence number N while the cache only virtual device 406 may correspond to the chunk of data 54 for sequence number N-1. Thus, when data is written by the host 22 to the local storage device 24, the data is provided to the cache 408 and an appropriate pointer of the cache only virtual device 404 is adjusted. Note that the data 10 will not be removed from the cache 408 until the data is destaged to the standard logical device 402 and the data is also released by the cache only virtual device 404, as described elsewhere herein.

In an embodiment herein, one of the cache only virtual devices 404, 406 is deemed “active” while the other is deemed “inactive”. Thus, for example, when the 15 sequence number N is even, the cache only virtual device 404 may be active while the cache only virtual device 406 is inactive. The active one of the cache only virtual devices 404, 406 handles writes from the host 22 while the inactive one of the cache only virtual devices 404, 406 corresponds to the data that is being transmitted from the local storage device 24 to the remote storage device 26.

While the data that is written by the host 22 is accumulated using the active one of the cache only virtual devices 404, 406 (for the sequence number N), the data corresponding to the inactive one of the cache only virtual devices 404, 406 (for previous sequence number N-1) is transmitted from the local storage device 24 to the remote storage device 26. For this and related embodiments, the DA's 35a-35c of the local storage device handle scanning the inactive one of the cache only virtual devices 404, 406 to send copy requests to one or more of the RA's 30a-30c to transmit the data from the local storage device 24 to the remote storage device 26. Thus, the steps 362, 374, discussed above in connection with suspending and resuming copying, may include providing messages/commands to the DA's 35a-35c.

Once the data has been transmitted to the remote storage device 26, the corresponding entry in the inactive one of the cache only virtual devices 404, 406 may be set to null. In addition, the data may also be removed from the cache 408 (i.e., the slot returned to the pool of slots for later use) if the data in the slot is not otherwise needed for another purpose (e.g., to be destaged to the standard logical device 402). A mechanism may be used to ensure that data is not removed from the cache 408 until all mirrors (including the cache only virtual devices 404, 406) are no longer using the data. Such a mechanism is described, for example, in U.S. Patent No. 5,537,568 issued on July 16, 1996 and in U.S. patent application no. 09/850,551 filed on July 7, 2001, both of which are incorporated by reference herein.

Referring to Figure 12, a flow chart 440 illustrates steps performed by the HA 28 in connection with a host 22 performing a write operation for embodiments where two COVD's are used by the R1 device to provide the system described herein. Processing begins at a first step 442 where a slot corresponding to the write is locked. In an 5 embodiment herein, each of the slots 412-414 of the cache 408 corresponds to a track of data on the standard logical device 402. Locking the slot at the step 442 prevents additional processes from operating on the relevant slot during the processing performed by the HA 28 corresponding to the steps of the flow chart 440.

Following the step 442 is a step 444 where a value for N, the sequence number, is 10 set. Just as with the embodiment that uses lists rather than COVD's on the R1 side, the value for the sequence number obtained at the step 444 is maintained during the entire write operation performed by the HA 28 while the slot is locked. As discussed elsewhere herein, the sequence number is assigned to each write to set the one of the chunks of data 52, 54 to which the write belongs. Writes performed by the host 22 are assigned the 15 current sequence number. It is useful that a single write operation maintain the same sequence number throughout.

Following the step 444 is a test step 446, which determines if the inactive one of the cache only virtual devices 404, 406 already points to the slot that was locked at the step 442 (the slot being operated upon). This may occur if a write to the same slot was 20 provided when the sequence number was one less than the current sequence number. The

data corresponding to the write for the previous sequence number may not yet have been transmitted to the remote storage device 26.

If it is determined at the test step 446 that the inactive one of the cache only virtual devices 404, 406 does not point to the slot, then control transfers from the test step 5 446 to another test step 448, where it is determined if the active one of the cache only virtual devices 404, 406 points to the slot. It is possible for the active one of the cache only virtual devices 404, 406 to point to the slot if there had been a previous write to the slot while the sequence number was the same as the current sequence number. If it is determined at the test step 448 that the active one of the cache only virtual devices 404, 10 406 does not point to the slot, then control transfers from the test step 448 to a step 452 where a new slot is obtained for the data. Following the step 452 is a step 454 where the active one of the cache only virtual devices 404, 406 is made to point to the slot.

Following the step 454, or following the step 448 if the active one of the cache only virtual devices 404, 406 points to the slot, is a step 456 where flags are set. At the 15 step 456, the RDF_WP flag (RDF write pending flag) is set to indicate that the slot needs to be transmitted to the remote storage device 26 using RDF. In addition, at the step 456, the IN_CACHE flag is set to indicate that the slot needs to be destaged to the standard logical device 402. Note that, in some instances, if the active one of the cache only virtual devices 404, 406 already points to the slot (as determined at the step 448) it is 20 possible that the RDF_WP and IN_CACHE flags were already set prior to execution of

the step 456. However, setting the flags at the step 456 ensures that the flags are set properly no matter what the previous state.

Following the step 456 is a step 458 where an indirect flag in the track table that points to the slot is cleared, indicating that the relevant data is provided in the slot and not 5 in a different slot indirectly pointed to. Following the step 458 is a step 462 where the data being written by the host 22 and the HA 28 is written to the slot. Following the step 462 is a step 464 where the slot is unlocked. Following step 464, processing is complete.

If it is determined at the test step 446 that the inactive one of the cache only virtual devices 404, 406 points to the slot, then control transfers from the step 446 to a 10 step 472, where a new slot is obtained. The new slot obtained at the step 472 may be used for the inactive one of the cache only virtual devices 404, 406 to effect the RDF transfer while the old slot may be associated with the active one of the cache only virtual devices 404, 406, as described below.

Following the step 472 is a step 474 where the data from the old slot is copied to 15 the new slot that was obtained at the step 472. Following the step 474 is a step 476 where the indirect flag (discussed above) is set to indicate that the track table entry for the inactive one of the cache only virtual devices 404, 406 points to the old slot but that the data is in the new slot which is pointed to by the old slot. Thus, setting indirect flag at the step 476 affects the track table of the inactive one of the cache only virtual devices 404, 20 406 to cause the track table entry to indicate that the data is in the new slot.

Following the step 476 is a step 478 where the mirror bits for the records in the new slot are adjusted. Any local mirror bits that were copied when the data was copied from the old slot to the new slot at the step 474 are cleared since the purpose of the new slot is to simply effect the RDF transfer for the inactive one of the cache only virtual devices. The old slot will be used to handle any local mirrors. Following the step 478 is the step 462 where the data is written to the slot. Following step 462 is the step 464 where the slot is unlocked. Following the step 464, processing is complete.

Referring to Figure 13, a flow chart 500 illustrates steps performed in connection with the local storage device 24 transmitting the chunk of data 54 to the remote storage device 26. The transmission essentially involves scanning the inactive one of the cache only virtual devices 404, 406 for tracks that have been written thereto during a previous iteration when the inactive one of the cache only virtual devices 404, 406 was active. In this embodiment, the DA's 35a-35c of the local storage device 24 scan the inactive one of the cache only virtual devices 404, 406 to copy the data for transmission to the remote storage device 26 by one or more of the RA's 30a-30c using the RDF protocol.

Processing begins at a first step 502 where the first track of the inactive one of the cache only virtual devices 404, 406 is pointed to in order to begin the process of iterating through all of the tracks. Following the first step 502 is a test step 504 where it is determined if the RDF_WP flag is set. As discussed elsewhere herein, the RDF_WP flag is used to indicate that a slot (track) contains data that needs to be transmitted via the RDF link. The RDF_WP flag being set indicates that at least some data for the slot

(track) is to be transmitted using RDF. In an embodiment herein, the entire slot is not transmitted. Rather, only records within the slot that have the appropriate mirror bits set (indicating the records have changed) are transmitted to the remote storage device 26. However, in other embodiments, it may be possible to transmit the entire slot, provided

5 that the remote storage device 26 only writes data corresponding to records having appropriate mirror bits set and ignores other data for the track, which may or may not be valid.

If it is determined at the test step 504 that the cache slot being processed has the RDF_WP flag set, then control transfers from the step 504 to a test step 505, where it is

10 determined if the slot contains the data or if the slot is an indirect slot that points to another slot that contains the relevant data. In some instances, a slot may not contain the data for the portion of the disk that corresponds to the slot. Instead, the slot may be an indirect slot that points to another slot that contains the data. If it is determined at the step 505 that the slot is an indirect slot, then control transfers from the step 505 to a step

15 506, where the data (from the slot pointed to by the indirect slot) is obtained. Thus, if the slot is a direct slot, the data for being sent by RDF is stored in the slot while if the slot is an indirect slot, the data for being sent by RDF is in another slot pointed to by the indirect slot.

Following the step 506 or the step 505 if the slot is a direct slot is a step 507

20 where data being sent (directly or indirectly from the slot) is copied by one of the DA's 35a-35c to be sent from the local storage device 24 to the remote storage device 26 using

the RDF protocol. Following the step 507 is a test step 508 where it is determined if the remote storage device 26 has acknowledged receipt of the data. If not, then control transfers from the step 508 back to the step 507 to resend the data. In other embodiments, different and more involved processing may be used to send data and

5 acknowledge receipt thereof. Such processing may include error reporting and alternative processing that is performed after a certain number of attempts to send the data have failed.

Once it is determined at the test step 508 that the data has been successfully sent, control passes from the step 508 to a step 512 to clear the RDF_WP flag (since the data

10 has been successfully sent via RDF). Following the step 512 is a step 514 where appropriate mirror flags are cleared to indicate that at least the RDF mirror (R2) no longer needs the data. In an embodiment herein, each record that is part of a slot (track) has individual mirror flags indicating which mirrors use the particular record. The R2 device is one of the mirrors for each of the records and it is the flags corresponding to the

15 R2 device that are cleared at the step 514.

Following the step 514 is a test step 516 which determines if any of the records of the track being processed have any other mirror flags set (for other mirror devices). If not, then control passes from the step 516 to a step 518 where the slot is released (i.e., no longer being used). In some embodiments, unused slots are maintained in a pool of slots

20 available for use. Note that if additional flags are still set for some of the records of the slot, it may mean that the records need to be destaged to the standard logical device 402

or are being used by some other mirror (including another R2 device). Following the step 518, or following the step 516 if more mirror flags are present, is a step 522 where the pointer that is used to iterate through each track entry of the inactive one of the cache only virtual devices 404, 406 is made to point to the next track. Following the step 522 is 5 a test step 524 which determines if there are more tracks of the inactive one of the cache only virtual devices 404, 406 to be processed. If not, then processing is complete. Otherwise, control transfers back to the test step 504, discussed above. Note that the step 522 is also reached from the test step 504 if it is determined that the RDF_WP flag is not set for the track being processed.

10 Referring to Figure 14, a diagram 700 illustrates a host 702 coupled to a plurality of local storage devices 703-705. The diagram 700 also shows a plurality of remote storage devices 706-708. Although only three local storage devices 703-705 and three remote storage devices 706-708 are shown in the diagram 700, the system described herein may be expanded to use any number of local and remote storage devices.

15 Each of the local storage devices 703-705 is coupled to a corresponding one of the remote storage devices 706-708 so that, for example, the local storage device 703 is coupled to the remote storage device 706, the local storage device 704 is coupled to the remote storage device 707 and the local storage device 705 is coupled to the remote storage device 708. The local storage device 703-705 and remote storage device 20 706-708 may be coupled using the ordered writes mechanism described herein so that, for example, the local storage device 703 may be coupled to the remote storage device 706

using the ordered writes mechanism. As discussed elsewhere herein, the ordered writes mechanism allows data recovery using the remote storage device in instances where the local storage device and/or host stops working and/or loses data.

In some instances, the host 702 may run a single application that simultaneously

5 uses more than one of the local storage devices 703-705. In such a case, the application may be configured to insure that application data is consistent (recoverable) at the local storage devices 703-705 if the host 702 were to cease working at any time and/or if one of the local storage devices 703-705 were to fail. However, since each of the ordered write connections between the local storage devices 703-705 and the remote storage

10 devices 706-708 is asynchronous from the other connections, then there is no assurance that data for the application will be consistent (and thus recoverable) at the remote storage devices 706-708. That is, for example, even though the data connection between the local storage device 703 and the remote storage device 706 (a first local/remote pair) is consistent and the data connection between the local storage device 704 and the remote

15 storage device 707 (a second local/remote pair) is consistent, it is not necessarily the case that the data on the remote storage devices 706, 707 is always consistent if there is no synchronization between the first and second local/remote pairs.

For applications on the host 702 that simultaneously use a plurality of local storage devices 703-705, it is desirable to have the data be consistent and recoverable at

20 the remote storage devices 706-708. This may be provided by a mechanism whereby the host 702 controls cycle switching at each of the local storage devices 703-705 so that the

data from the application running on the host 702 is consistent and recoverable at the remote storage devices 706-708. This functionality is provided by a special application that runs on the host 702 that switches a plurality of the local storage devices 703-705 into multi-box mode, as described in more detail below.

- 5 Referring to Figure 15, a table 730 has a plurality of entries 732-734. Each of the entries 732-734 correspond to a single local/remote pair of storage devices so that, for example, the entry 732 may correspond to pair of the local storage device 703 and the remote storage device 706, the entry 733 may correspond to pair of the local storage device 704 and the remote storage device 707 and the entry 734 may correspond to the 10 pair of local storage device 705 and the remote storage device 708. Each of the entries 732-734 has a plurality of fields where a first field 736a-736c represents a serial number of the corresponding local storage device, a second field 738a-738c represents a session number used by the multi-box group, a third field 742a-742c represents the serial number of the corresponding remote storage device of the local/remote pair, and a fourth field 15 744a-744c represents the session number for the multi-box group. The table 730 is constructed and maintained by the host 702 in connection with operating in multi-box mode. In addition, the table 730 is propagated to each of the local storage devices and the remote storage devices that are part of the multi-box group. The table 730 may be used to facilitate recovery, as discussed in more detail below.
- 20 Different local/remote pairs may enter and exit multi-box mode independently in any sequence and at any time. The host 702 manages entry and exit of local storage

device/remote storage device pairs into and out of multi-box mode. This is described in more detail below.

Referring to Figure 16, a flowchart 750 illustrates steps performed by the host 702 in connection with entry or exit of a local/remote pair in to or out of multi-box mode.

- 5 Processing begins at a first step 752 where multi-box mode operation is temporarily suspended. Temporarily suspending multi-box operation at the step 752 is useful to facilitate the changes that are made in connection with entry or exit of a remote/local pair in to or out of multi-box mode. Following the step 752, is a step 754 where a table like the table 730 of Figure 15 is modified to either add or delete an entry, as appropriate.
- 10 Following the step 754 is a step 756 where the modified table is propagated to the local storage devices and remote storage devices of the multi-box group. Propagating the table at the step 756 facilitates recovery, as discussed in more detail elsewhere herein.

Following the step 756 is a step 758 where a message is sent to the affected local storage device to provide the change. The local storage device may configure itself to

- 15 run in multi-box mode or not, as described in more detail elsewhere herein. As discussed in more detail below, a local storage device handling ordered writes operates differently depending upon whether it is operating as part of a multi-box group or not. If the local storage device is being added to a multi-box group, the message sent at the step 758 indicates to the local storage device that it is being added to a multi-box group so that the
- 20 local storage device should configure itself to run in multi-box mode. Alternatively, if a local storage device is being removed from a multi-box group, the message sent at the

step 758 indicates to the local storage device that it is being removed from the multi-box group so that the local storage device should configure itself to not run in multi-box mode.

- Following step 758 is a test step 762 where it is determined if a local/remote pair
- 5 is being added to the multi-box group (as opposed to being removed). If so, then control transfers from the test step 762 to a step 764 where tag values are sent to the local storage device that is being added. The tag values are provided with the data transmitted from the local storage device to the remote storage device in a manner similar to providing the sequence numbers with the data. The tag values are controlled by the host and set so that
- 10 all of the local/remote pairs send data having the same tag value during the same cycle.

Use of the tag values is discussed in more detail below. Following the step 764, or following the step 762 if a new local/remote pair is not being added, is a step 766 where multi-box operation is resumed. Following the step 766, processing is complete.

- Referring to Figure 17, a flow chart 780 illustrates steps performed in connection
- 15 with the host managing cycle switching for multiple local/remote pairs running as a group in multi-box mode. As discussed elsewhere herein, multi-box mode involves having the host synchronize cycle switches for more than one remote/local pair to maintain data consistency among the remote storage devices. Cycle switching is coordinated by the host rather than being generated internally by the local storage
- 20 devices. This is discussed in more detail below.

Processing for the flow chart 780 begins at a test step 782 which determines if M seconds have passed. Just as with non-multi-box operation, cycle switches occur no sooner than every M seconds where M is a number chosen to optimize various performance parameters. As the number M is increased, the amount of overhead 5 associated with switching decreases. However, increasing M also causes the amount of data that may be potentially lost in connection with a failure to also increase. In an embodiment herein, M is chosen to be thirty seconds, although, obviously other values for M may be used.

- If it is determined at the test step 782 that M seconds have not passed, then 10 control transfers back to the step 782 to continue waiting until M seconds have passed. Once it is determined at the test step 782 that M seconds have passed, control transfers from the step 782 to a step 784 where the host queries all of the local storage devices in the multi-box group to determine if all of the local/remote pairs are ready to switch. The local/remote pairs being ready to switch is discussed in more detail hereinafter.
- 15 Following the step 784 is a test step 786 which determines if all of the local/remote pairs are ready to switch. If not, control transfers back to the step 784 to resume the query. In an embodiment herein, it is only necessary to query local/remote pairs that were previously not ready to switch since, once a local/remote pair is ready to switch, the pair remains so until the switch occurs.

Once it is determined at the test step 786 that all of the local/remote pairs in the multi-box group are ready to switch, control transfers from the step 786 to a step 788 where an index variable, N, is set equal to one. The index variable N is used to iterate through all the local/remote pairs (i.e., all of the entries 732-734 of the table 730 of 5 Figure 15). Following the step 788 is a test step 792 which determines if the index variable, N, is greater than the number of local/remote pairs in the multi-box group. If not, then control transfers from the step 792 to a step 794 where an open window is performed for the Nth local storage device of the Nth pair by the host sending a command (e.g., an appropriate system command) to the Nth local storage device. Opening the 10 window for the Nth local storage device at the step 794 causes the Nth local storage device to suspend writes so that any write by a host that is not begun prior to opening the window at the step 794 will not be completed until the window is closed (described below). Not completing a write operation prevents a second dependant write from occurring prior to completion of the cycle switch. Any writes in progress that were 15 begun before opening the window may complete prior to the window being closed.

Following the step 794 is a step 796 where a cycle switch is performed for the Nth local storage device. Performing the cycle switch at the step 796 involves sending a command from the host 702 to the Nth local storage device. Processing the command from the host by the Nth local storage device is discussed in more detail below. Part of 20 the processing performed at the step 796 may include having the host provide new values for the tags that are assigned to the data. The tags are discussed in more detail elsewhere herein. In an alternative embodiment, the operations performed at the steps 794, 796 may

be performed as a single integrated step 797, which is illustrated by the box drawn around the steps 794, 796.

Following the step 796 is a step 798 where the index variable, N, is incremented.

Following step 798, control transfers back to the test step 792 to determine if the index

5 variable, N, is greater than the number of local/remote pairs.

If it is determined at the test step 792 that the index variable, N, is greater than the number of local/remote pairs, then control transfers from the test step 792 to a step 802

where the index variable, N, is set equal to one. Following the step 802 is a test step 804

which determines if the index variable, N, is greater than the number of local/remote

10 pairs. If not, then control transfers from the step 804 to a step 806 where the window for the Nth local storage device is closed. Closing the window of the step 806 is performed by the host sending a command to the Nth local storage device to cause the Nth local

storage device to resume write operations. Thus, any writes in process that were

suspended by opening the window at the step 794 may now be completed after execution

15 of the step 806. Following the step 806, control transfers to a step 808 where the index variable, N, is incremented. Following the step 808, control transfers back to the test step 804 to determine if the index variable, N, is greater than the number of local/remote pairs.

If so, then control transfers from the test step 804 back to the step 782 to begin processing

for the next cycle switch.

Referring to Figure 18, a flow chart 830 illustrates steps performed by a local storage device in connection with cycle switching. The flow chart 830 of Figure 18 replaces the flow chart 350 of Figure 10 in instances where the local storage device supports both multi-box mode and non-multi-box mode. That is, the flow chart 830

5 shows steps performed like those of the flow chart 350 of Figure 10 to support non-multi-box mode and, in addition, includes steps for supporting multi-box mode.

Processing begins at a first test step 832 which determines if the local storage device is operating in multi-box mode. Note that the flow chart 750 of Figure 16 shows the step 758 where the host sends a message to the local storage device. The message sent at the step 758 indicates to the local storage device whether the local storage device is in multi-box mode or not. Upon receipt of the message sent by the host at the step 758, the local storage device sets an internal variable to indicate whether the local storage device is operating in multi-box mode or not. The internal variable may be examined at the test step 832.

10 If it is determined at the test step 832 that the local storage device is not in multi-box mode, then control transfers from the test step 832 to a step 834 to wait M seconds for the cycle switch. If the local storage device is not operating in multi-box mode, then the local storage device controls its own cycle switching and thus executes the step 834 to wait M seconds before initiating the next cycle switch.

Following the step 834, or following the step 832 if the local storage device is in multi-box mode, is a test step 836 which determines if all of the HA's of the local storage device have set a bit indicating that the HA's have completed all of the I/O's for a previous sequence number. When the sequence number changes, each of the HA's

5 notices the change and sets a bit indicating that all I/O's of the previous sequence number are completed. For example, if the sequence number changes from N-1 to N, an HA will set the bit when the HA has completed all I/O's for sequence number N-1. Note that, in some instances, a single I/O for an HA may take a long time and may still be in progress even after the sequence number has changed. Note also that, for some systems, a

10 different mechanism may be used to determine if all HA's have completed their N-1 I/O's. The different mechanism may include examining device tables. Once it is determined at the test step 836 that all HA's have set the appropriate bit, control transfers from the test step 836 to a step 888 which determines if the inactive chunk for the local storage device is empty. Once it is determined at the test step 888 that the inactive chunk

15 is empty, control transfers from the step 888 to a step 899, where copying of data from the local storage device to the remote storage device is suspended. It is useful to suspend copying data until the sequence number switch is complete.

Following the step 899 is a test step 892 to determine if the local storage device is in multi-box mode. If it is determined at the test step 892 that the local storage device is

20 in multi-box mode, then control transfers from the test step 892 to a test step 894 to determine if the active chunk of the corresponding remote storage device is empty. As discussed in more detail below, the remote storage device sends a message to the local

storage device once it has emptied its active chunk. In response to the message, the local storage device sets an internal variable that is examined at the test step 894.

Once it is determined at the test step 894 that the active chunk of the remote storage device is empty, control transfers from the test step 894 to a step 896 where an 5 internal variable is set on a local storage device indicating that the local storage device is ready to switch cycles. As discussed above in connection with the flow chart 780 of Figure 17, the host queries each of the local storage devices to determine if each of the local storage devices are ready to switch. In response to the query provided by the host, the local storage device examines the internal variable set at the step 896 and returns the 10 result to the host.

Following step 896 is a test step 898 where the local storage device waits to receive the command from the host to perform the cycle switch. As discussed above in connection with the flow chart 780 of Figure 17, the host provides a command to switch cycles to the local storage device when the local storage device is operating in multi-box mode. Thus, the local storage device waits for the command at the step 898, which is 15 only reached when the local storage device is operating in multi-box mode.

Once the local storage device has received the switch command from the host, control transfers from the step 898 to a step 902 to send a commit message to the remote storage device. Note that the step 902 is also reached from the test step 892 if it is 20 determined at the step test 892 that the local storage device is not in multi-box mode. At

the step 902, the local storage device sends a commit message to the remote storage device. In response to receiving a commit message for a particular sequence number, the remote storage device will begin restoring the data corresponding to the sequence number, as discussed above.

5 Following the step 902 is a step 906 where the sequence number is incremented and a new value for the tag (from the host) is stored. The sequence number is as discussed above. The tag is the tag provided to the local storage device at the step 764 and at the step 796, as discussed above. The tag is used to facilitate data recovery, as discussed elsewhere herein.

10 Following the step 906 is a step 907 where completion of the cycle switch is confirmed from the local storage device to the host by sending a message from the local storage device to the host. In some embodiments, it is possible to condition performing the step 907 on whether the local storage device is in multi-box mode or not, since, if the local storage device is not in multi-box mode, the host is not necessarily interested in
15 when cycle switches occur.

Following the step 907 is a step 908 where the bits for the HA's that are used in the test step 836 are all cleared so that the bits may be set again in connection with the increment of the sequence number. Following the step 908 is a test step 912 which determines if the remote storage device has acknowledged the commit message. Note
20 that if the local/remote pair is operating in multi-box mode and the remote storage device

active chunk was determined to be empty at the step 894, then the remote storage device should acknowledge the commit message nearly immediately since the remote storage device will be ready for the cycle switch immediately because the active chunk thereof is already empty.

5 Once it is determined at the test step 912 that the commit message has been acknowledged by the remote storage device, control transfers from the step 912 to a step 914 where the suspension of copying, which was provided at the step 899, is cleared so that copying from the local storage device to the remote storage device may resume.

Following the step 914, processing is complete.

10 Referring to Figure 19, a flow chart 940 illustrates steps performed in connection with RA's scanning the inactive buffers to transmit RDF data from the local storage device to the remote storage device. The flow chart 940 of Figure 19 is similar to the flow chart 200 of Figure 6 and similar steps are given the same reference number. However, the flow chart 940 includes two additional steps 942, 944 which are not found

15 in the flow chart 200 of Figure 6. The additional steps 942, 944 are used to facilitate multi-box processing. After data has been sent at the step 212, control transfers from the step 212 to a test step 942 which determines if the data being sent is the last data in the inactive chunk of the local storage device. If not, then control transfers from the step 942 to the step 214 and processing continues as discussed above in connection with the flow

20 chart 200 of Figure 6. Otherwise, if it is determined at the test step 942 that the data being sent is the last data of the chunk, then control transfers from the step 942 to the step

944 to send a special message from the local storage device to the remote storage device indicating that the last data has been sent. Following the step 944, control transfers to the step 214 and processing continues as discussed above in connection with the flow chart 200 of Figure 6. In some embodiments, the steps 942, 944 may be performed by a 5 separate process (and/or separate hardware device) that is different from the process and/or hardware device that transfers the data.

Referring to Figure 20, a flow chart 950 illustrates steps performed in connection with RA's scanning the inactive buffers to transmit RDF data from the local storage device to the remote storage device. The flow chart 950 of Figure 20 is similar to the 10 flow chart 500 of Figure 13 and similar steps are given the same reference number. However, the flow chart 950 includes an additional step 952, which is not found in the flow chart 500 of Figure 13. The additional steps 952 is used to facilitate multi-box processing and is like the additional step 944 of the flowchart 940 of Figure 19. After it is determined at the test step 524 that no more slots remain to be sent from the local 15 storage device to the remote storage device, control transfers from the step 524 to the step 952 to send a special message from the local storage device to the remote storage device indicating that the last data for the chunk has been sent. Following the step 952, processing is complete.

Referring to Figure 21, a flow chart 960 illustrates steps performed at the remote 20 storage device in connection with providing an indication that the active chunk of the remote storage device is empty. The flow chart 960 is like the flow chart 300 of Figure 9

except that the flow chart 960 shows a new step 962 that is performed after the active chunk of the remote storage device has been restored. At the step 962, the remote storage device sends a message to the local storage device indicating that the active chunk of the remote storage device is empty. Upon receipt of the message sent at the step 962, the
5 local storage device sets an internal variable indicating that the inactive buffer of the remote storage device is empty. The local variable is examined in connection with the test step 894 of the flow chart 830 of Figure 18, discussed above.

Referring to Figure 22, a diagram 980 illustrates the host 702, local storage devices 703-705 and remote storage devices 706-708, that are shown in the diagram 700
10 of Figure 14. The Diagram 980 also includes a first alternative host 982 that is coupled to the host 702 and the local storage devices 703-705. The diagram 980 also includes a second alternative host 984 that is coupled to the remote storage devices 706-708. The alternative hosts 982, 984 may be used for data recovery, as described in more detail below.
15

When recovery of data at the remote site is necessary, the recovery may be performed by the host 702 or, by the host 982 provided that the links between the local storage devices 703-705 and the remote storage devices 706-708 are still operational. If the links are not operational, then data recovery may be performed by the second alternative host 984 that is coupled to the remote storage devices 706-708. The second
20 alternative host 984 may be provided in the same location as one or more of the remote storage devices 706-708. Alternatively, the second alternative host 984 may be remote

from all of the remote storage devices 706-708. The table 730 that is propagated throughout the system is accessed in connection with data recovery to determine the members of the multi-box group.

Referring to Figure 23, a flow chart 1000 illustrates steps performed by each of

5 the remote storage devices 706-708 in connection with the data recovery operation. The steps of the flowchart 1000 may be executed by each of the remote storage devices 706-708 upon receipt of a signal or a message indicating that data recovery is necessary. In some embodiments, it may be possible for a remote storage device to automatically sense that data recovery is necessary using, for example, conventional criteria such as length of

10 time since last write.

Processing begins at a first step 1002 where the remote storage device finishes restoring the active chunk in a manner discussed elsewhere herein. Following the step 1002 is a test step 1004 which determines if the inactive chunk of the remote storage device is complete (i.e., all of the data has been written thereto). Note that a remote

15 storage device may determine if the inactive chunk is complete using the message sent by the local storage device at the steps 944, 952, discussed above. That is, if the local storage device has sent the message at the step 944 or the step 952, then the remote storage device may use receipt of that message to confirm that the inactive chunk is complete.

If it is determined at the test step 1004 that the inactive chunk of the remote storage device is not complete, then control transfers from the test step 1004 to a step 1006 where the data from the inactive chunk is discarded. No data recovery is performed using incomplete inactive chunks since the data therein may be inconsistent with the 5 corresponding active chunks. Accordingly, data recovery is performed using active chunks and, in some cases, inactive chunks that are complete. Following the step 1006, processing is complete.

If it is determined at the test step 1004 that the inactive chunk is complete, then control transfers from the step 1004 to the step 1008 where the remote storage device 10 waits for intervention by the host. If an inactive chunk, one of the hosts 702, 982, 984, as appropriate, needs to examine the state of all of the remote storage devices in the multi-box group to determine how to perform the recovery. This is discussed in more detail below.

Following step 1008 is a test step 1012 where it is determined if the host has 15 provided a command to all storage device to discard the inactive chunk. If so, then control transfers from the step 1012 to the step 1006 to discard the inactive chunk. Following the step 1006, processing is complete.

If it is determined at the test step 1002 that the host has provided a command to restore the complete inactive chunk, then control transfers from the step 1012 to a step 20 1014 where the inactive chunk is restored to the remote storage device. Restoring the

inactive chunk in the remote storage device involves making the inactive chunk an active chunk and then writing the active chunk to the disk as described elsewhere herein.

Following the step 1014, processing is complete.

Referring to Figure 24, a flow chart 1030 illustrates steps performed in connection
5 with one of the hosts 702, 982, 984 determining whether to discard or restore each of the
inactive chunks of each of the remote storage devices. The one of the hosts 702, 982, 984
that is performing the restoration communicates with the remote storage devices 706-708
to provide commands thereto and to receive information therefrom using the tags that are
assigned by the host as discussed elsewhere herein.

10 Processing begins at a first step 1032 where it is determined if any of the remote
storage devices have a complete inactive chunk. If not, then there is no further
processing to be performed and, as discussed above, the remote storage devices will
discard the incomplete chunks on their own without host intervention. Otherwise, control
transfers from the test step 1032 to a test step 1034 where the host determines if all of the
15 remote storage devices have complete inactive chunks. If so, then control transfers from
the test step 1034 to a test step 1036 where it is determined if all of the complete inactive
chunks of all of the remote storage devices have the same tag number. As discussed
elsewhere herein, tags are assigned by the host and used by the system to identify data in
a manner similar to the sequence number except that tags are controlled by the host to
20 have the same value for the same cycle.

If it is determined at the test step 1036 that all of the remote storage devices have the same tag for the inactive chunks, then control transfers from the step 1036 to a step 1038 where all of the inactive chunks are restored. Performing the step 1038 ensures that all of the remote storage devices have data from the same cycle. Following the step 5 1038, processing is complete.

If it is determined at the test step 1034 that all of the inactive chunks are not complete, or if it is determined that at the step 1036 that all of the complete inactive chunks do not have the same tag, then control transfers to a step 1042 where the host provides a command to the remote storage devices to restore the complete inactive 10 chunks having the lower tag number. For purposes of explanation, it is assumed that the tag numbers are incremented so that a lower tag number represents older data. By way of example, if a first remote storage device had a complete inactive chunk with a tag value of three and a second remote storage device had a complete inactive chunk with a tag value of four, the step 1042 would cause the first remote storage device (but not the 15 second) to restore its inactive chunk. Following the step 1042 is a step 1044 where the host provides commands to the remote storage devices to discard the complete inactive buffers having a higher tag number (e.g., the second remote storage device in the previous example). Following step 1044, processing is complete.

Following execution of the step 1044, each of the remote storage devices contains 20 data associated with the same tag value as data for the other ones of the remote storage

devices. Accordingly, the recovered data on the remote storage devices 706-708 should be consistent.

While the invention has been disclosed in connection with various embodiments, modifications thereon will be readily apparent to those skilled in the art. Accordingly,
5 the spirit and scope of the invention is set forth in the following claims.